

AD-A050 679

CONTROL DATA CORP MINNEAPOLIS MINN DIGITAL IMAGE SYS--ETC F/6 14/5
IMAGE COMPRESSION TECHNIQUES.(U)

DEC 77 A E LABONTE, C J MCCALLUM

F30602-76-C-0350

UNCLASSIFIED

RADC-TR-77-405

NL

1 of 2
AD
A050679



AD A 050679

DDC FILE COPY


RADC-TR-77-405
Final Technical Report
December 1977



IMAGE COMPRESSION TECHNIQUES

Dr. A.E. Labonte
Ms. C.J. McCallum

Control Data Corporation



Approved for public release; distribution unlimited.

ROME AIR DEVELOPMENT CENTER
Air Force Systems Command
Griffiss Air Force Base, New York 13441

This report has been reviewed by the RADC Information Office (OI) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RADC-TR-77-405 has been reviewed and is approved for publication.

APPROVED:

Peter J. Costianes

PETER J. COSTIANES
Project Engineer

APPROVED:

Howard Davis

HOWARD DAVIS
Technical Director
Intelligence & Reconnaissance Division

FOR THE COMMANDER:

John P. Huss

JOHN P. HUSS
Acting Chief, Plans Office

If your address has changed or if you wish to be removed from the RADC mailing list, or if the addressee is no longer employed by your organization, please notify RADC (IRRE) Griffiss AFB NY 13441. This will assist us in maintaining a current mailing list.

Do not return this copy. Retain or destroy.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

19 REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
18 RADC-TR-77-405		
4. TITLE (and Subtitle)		5. TYPE OF REPORT & PERIOD COVERED
6 IMAGE COMPRESSION TECHNIQUES.		9 Final Technical Report, 30 Jun 77 - 30 Jun 78.
7. AUTHOR(s)		6. PERFORMING ORG. REPORT NUMBER
Dr. A.E. Labonte Mr. C.J. McCallum		N/A
9. PERFORMING ORGANIZATION NAME AND ADDRESS		8. CONTRACT OR GRANT NUMBER(s)
Control Data Corporation/Digital Image Systems Div. 2800 East Old Shakopee Road Minneapolis MN 55420		15 F30602-76-C-0350
11. CONTROLLING OFFICE NAME AND ADDRESS		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBER
Rome Air Development Center (IRRE) ✓ Griffiss AFB NY 13441		P.E. 62702F J.O. 62441268 16 6244 17 12
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE
Same		11 Dec 1977
		13. NUMBER OF PAGES
		93
		15. SECURITY CLASS. (of this report)
		UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
		N/A
16. DISTRIBUTION STATEMENT (of this Report)		
Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
Same		
18. SUPPLEMENTARY NOTES		
RADC Project Engineer: Peter J. Costianes (IRRE)		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
Intelligence Optical Detection Optics		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)		
Timely transmission of large format digital imagery over narrow bandwidth lines requires efficient and high compression of the digitized images. The techniques developed, Micro-Adaptive Picture Sequencing (MAPS), is a two-dimensional, spatial adaptive technique which uses the Redundant Area Coding (REARC) concept along with a very versatile algorithm developed by Control Data Corporation. Compression ratios of 30:1 have been achieved with MSE ranging from .548 to 2.534 percent for a broad variety of visible and radar imagery.		

DDC
RECEIVED
MAR 3 1978
F

DD FORM 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

408 732

JOB

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE/When Data Entered

[Faint, mostly illegible text and markings within a large rectangular frame, possibly a redacted document or a form with very light print.]

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE/When Data Entered

PREFACE

The following document constitutes the technical final report on Contract F30602-76-C-0350, IMAGE COMPRESSION TECHNIQUES. The work was performed in the Digital Image Systems Division of Control Data Corporation by Dr. A. E. LaBonte, principal investigator, and Ms. C. J. McCallum, digital processing.

Thanks are extended to Mr. P. Costianes of the Rome Air Development Center for his continuing enthusiasm and technical guidance throughout the program.

ACCESSION for	
NTIS	W. C. 300 ON <input checked="" type="checkbox"/>
DDC	B. H. Section <input type="checkbox"/>
UNANNOUNCED	<input type="checkbox"/>
JUS TICATION	
PY	
DISTRIBUTION/AVAILABILITY CODES	
01	SPECIAL
A	

TABLE OF CONTENTS

<u>Section</u>	<u>Title</u>	<u>Page</u>
1	SUMMARY	1-1
	1.1 Criteria	1-1
	1.2 Technique Selected - Principal Conclusion	1-1
	1.3 Organization of the Report	1-2
2	COMPRESSION TECHNIQUES	2-1
	2.1 Target/Background Partition - The REARC Technique .	2-1
	2.2 Algorithm Classification	2-2
	2.2.1 Entropy Encoding	2-2
	2.2.2 Spatial Encoding	2-4
	2.2.3 Transform Encoding	2-5
	2.3 Conclusion	2-6
3	TWO-DIMENSIONAL IMAGE CODING BY MICRO-ADAPTIVE PICTURE SEQUENCING	3-1
	3.1 Micro-Adaptive Picture Sequencing	3-3
	3.1.1 Compression	3-3
	3.1.2 Decompression	3-11
	3.1.3 Interactive Macro-Fidelity Control	3-15
	3.2 MAPS Examples	3-16
	3.2.1 IEEE Facsimile Test Chart	3-16
	3.2.2 Air Field Scene	3-20
	3.2.3 Power Plant Scene	3-20
4	INTRA-MAPS TRADE-OFF STUDIES	4-1
	4.1 Composite Test Image	4-1
	4.2 Selection of MAPS Contrast Control Mode	4-4
	4.3 Selection of MAPS Highest Resolution Code	4-5
	4.4 Selection of MAPS Maximum Block Size	4-11
	4.5 Summary of Intra-MAPS Trade-Off Selections	4-14
5	TEST IMAGERY RESULTS	5-1
	5.1 Tabulation of MAPS Performance	5-1
	5.2 MAPS Imagery	5-2

TABLE OF CONTENTS (Cont.)

<u>Section</u>	<u>Title</u>	<u>Page</u>
	5.3 Line Error Effects	5-36
	5.3.1 Line Error Model	5-37
	5.3.2 MAPS Line Error Example	5-40
6	IMPLEMENTATION RECOMMENDATIONS	6-1
7	MAPS EVOLUTION	7-1
	7.1 Adaptive Decompression	7-1
	7.2 Control Matrix Selection	7-2
	7.3 Error Recovery Optimization	7-2
	7.4 Interactive Macro-Fidelity Control	7-2
	7.5 On-Line MAPS Demonstration	7-3
	References	R-1

LIST OF FIGURES

<u>Figure</u>	<u>Title</u>	<u>Page</u>
2-1	Target/Background Compression Requirements.....	2-3
2-2	Compression Algorithm Classification Net.....	2-8
3-1	The Bases of MAPS.....	3-2
	(a) Image Description	
	(b) Vision Heuristic	
3-2	Original Image Subframe Partition.....	3-5
	(a) Local Position Numbering (all levels $0 \leq L \leq M$)	
	(b) Level Nesting, $M = 3$	
	(c) 'Block' Sequence Path (Original Image)	
3-3	Sample MAPS Subframe Partition.....	3-7
	(a) A Possible MAPS Configuration	
	(b) An Excluded Configuration	
	(c) 'Block' Sequence Path	
3-4	Micro-Fidelity Control.....	3-8
	(a) Contrast Definition	
	(b) Contrast Control Matrix	
	(c) Patterns Emphasized by a Small Middle Step Threshold	
3-5	MAPS Subframe Block Compression - Functional Flow.....	3-12
3-6	Adaptive Decompression Geometry.....	3-14
3-7	IEEE Facsimile Test Chart.....	3-17
	(a) Original	
	(b) Block Mode 0.593 bits/pixel MSE = 0.822	
	(c) Adaptive Mode 0.593 bits/pixel MSE = 0.691	
3-8	Airfield Scene.....	3-18
	(a) Original (Spotlighted)	
	(b) Fixed Fidelity 0.455 bits/pixel	
	(c) Variable Fidelity 0.198 bits/pixel	

LIST OF FIGURES (Cont.)

<u>Figure</u>	<u>Title</u>	<u>Page</u>
3-9	Power Plant Scene.....	3-19
	(a) Power Plant Scene (Spotlighted Region) 6-bit Original	
	(b) 0.167 bits/pixel overall, 0.762 bits/pixel spotlight, 0.119 bits/pixel background	
4-1	Composite Test Image - Original.....	4-2
4-2	Fidelity vs Compression - 48 Subframe Composite Image.....	4-3
4-3	Extreme Contrast Control Mode	
	(a) Compression/MSE/Operations Map.....	4-6
	(b) 48 Subframe Composite.....	4-7
4-4	Step Contrast Control Mode	
	(a) Compression/MSE/Operations Map.....	4-8
	(b) 48 Subframe Composite.....	4-9
4-5	Comparative Compressions for LMAX = 3 and LMAX = 4.....	4-12
4-6	Intra-MAPS Trade-off Selections.....	4-15
5-1	(a) Aerial Photo - Targeted Original.....	5-6
	(b) Aerial Photo - 14.46:1.....	5-7
	(c) Aerial Photo - 21.84:1.....	5-8
	(d) Aerial Photo - 30.29:1.....	5-9
5-2	(a) Riverfront - Targeted Original.....	5-10
	(b) Riverfront - 30.02:1.....	5-11
5-3	(a) Radar Scene - Targeted Original.....	5-12
	(b) Radar Scene - 20.55:1.....	5-13
	(c) Radar Scene - 30.17:1.....	5-14
	(d) Radar Scene - 38.95:1.....	5-15
5-4	(a) Harbor 8 x 8 - Targeted Original.....	5-16
	(b) Harbor 8 x 8 - 30.78:1.....	5-17
5-5	(a) Harbor 4 x 4 - Targeted Original.....	5-18
	(b) Harbor 4 x 4 - 22.62:1.....	5-19
	(c) Harbor 4 x 4 - 30.03:1.....	5-20
	(d) Harbor 4 x 4 - 38.28:1.....	5-21

LIST OF FIGURES (CONT).

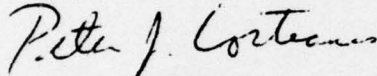
<u>Figure</u>	<u>Title</u>	<u>Page</u>
5-6	(a) Harbor 2 x 2 - Targeted Original.....	5-22
	(b) Harbor 2 x 2 - 30.03:1.....	5-23
5-7	(a) IEEE Chart - Original.....	5-24
	(b) IEEE Chart - 10.12:1.....	5-25
	(c) IEEE Chart - 16.26:1.....	5-26
	(d) IEEE Chart - 30.76:1.....	5-27
5-8	(a) SAM Site - Targeted Original.....	5-28
	(b) SAM Site - 30.10:1.....	5-29
5-9	(a) Airfield - Targeted Original.....	5-30
	(b) Airfield - 21.83:1.....	5-31
	(c) Airfield - 30.34:1.....	5-32
	(d) Airfield - 35.34:1.....	5-33
5-10	(a) IEEE Girl - Original.....	5-34
	(b) IEEE Girl - 30.54:1.....	5-35
5-11	Line Error Model Relations.....	5-38
5-12	Mean Error Rate vs α for the Truncated Pareto Distribution...	5-39
5-13	Line Error Intensity Effects: Airfield - 30.34:1.....	5-41

LIST OF TABLES

<u>Table</u>	<u>Title</u>	<u>Page</u>
5-1	FULL FRAME COMPRESSIONS - COMPRESSION RATIO.....	5-3
5-2	FULL FRAME COMPRESSIONS - MEAN SQUARE ERROR.....	5-4
5-3	FULL FRAME COMPRESSIONS - OPERATIONS COUNT.....	5-5

EVALUATION

This effort has succeeded in producing an image compression technique which is capable of 30:1 image compression with reasonable fidelity and ease of implementation. This will enable rapid transmission of imagery over existing narrow-band communication links which represents a capability that is urgently needed by many users of photo intelligence in the DoD. This is consistent with the objectives of TPO R2C which establishes a need for a capability to securely transmit digital image data with improved compression ratios over existing communication assets.



PETER J. COSTIANES
Project Engineer

SECTION ONE

SUMMARY

Timely transmission of large-format digital imagery over narrow bandwidth land lines requires efficient and high compression of the digitized images. The problem addressed in this effort consists of searching, selecting, and optimizing among various image compression techniques to extend the state-of-the-art to compression ratios of 30:1 or higher.

1.1 CRITERIA

Criteria for technique selection involve:

- Output image fidelity,
- Achievable compression ratio (at least 30:1),
- Susceptibility to transmission bit errors, and
- Implementation cost and speed.

In addition, the selected technique must exploit the variable fidelity requirements imposed by the REDundant AREA Coding (REARC) scheme developed previously through RADC. With REARC, the image is coarsely partitioned into 'target' and 'background' regions such that the background is used only for context and may thus be retained at considerably reduced fidelity.

1.2 TECHNIQUE SELECTED - PRINCIPAL CONCLUSION

Selection of Micro-Adaptive Picture Sequencing (MAPS), a two-dimensional image compression technique newly developed by the Digital Image Systems Division of Control Data Corporation, constitutes the central result of this effort.

MAPS has dramatic implementation advantages over the few other available techniques (all in the adaptive transform classification) which have the potential to operate in this high compression regime. Furthermore, the MAPS

concepts are very robust, promising extensive and diverse areas for enhancement and evolution of the technique. The body of the report presenting the details of these assertions is outlined in the following subsection.

1.3 ORGANIZATION OF THE REPORT

Section Two of the text evaluates the general implications of the high compression goals. This analysis establishes the need to use a destructive technique - i.e. to allow some degradation - for both the background and target regions. This is followed by a brief survey which classifies available image compression algorithms and places MAPS in this context.

Section Three provides a stand-alone description of Micro-Adaptive Picture Sequencing since it has not at this point been documented in the open technical literature.* This section establishes the MAPS conceptual framework and gives a procedural description sufficient for basic MAPS implementation.

Section Four explores three MAPS trade-off options in terms of five performance dimensions - compression, fidelity, error recovery, operations count, and memory requirements. Special attention is directed to preservation of MAPS inherent error recovery potential and to the selection of a minimum memory configuration consistent with the compression goals.

Section Five presents the results of applying MAPS to a test set of ten diverse full frame images; each original consisting of 2048 x 2048 6-bit pixels. An example with compression ratio just over 30:1 is given for each original and five of the scenes are presented at two more compression levels. The section concludes with an example of the effects of transmission line errors on the output image intensities; a truncated Pareto distribution is used to model the error statistics.

* Essentially this same material will appear in the Proceedings of the Society of Photo-Optical Instrumentation Engineers, Volume 119, 1977.

Section Six recommends incorporation of MAPS in the microcode of an existing machine - the CDC[®] Flexible Processor - should immediate implementation be desired. This approach would retain the capability to respond to the expected evolution of MAPS through modification of the corresponding software.

Section Seven concludes the report with an outline of several near-term developments which are recommended for further evolution within the framework of Micro-Adaptive Picture Sequencing.

SECTION TWO

COMPRESSION TECHNIQUES

In order to enable image transmission over narrow bandwidth land lines, a goal of very high compression ratio has been established - in excess of 30:1 overall for a typical 2048 x 2048 image. Achievement of this goal is helped by partitioning the image into target and background regions using the Redundant Area Coding (REARC)¹ technique and relaxing the fidelity criteria required in the background portions of the image. Nevertheless, the very high overall compression implies that destructive coding must be applied to both target and background.

Review of available coding techniques reveals that they may be classified into three broad types - entropy, spatial, and transform coding. Among these, only the latter two exhibit the potential for reaching the required compressions. A further brief comparison selects a new spatial technique, Micro-Adaptive Picture Sequencing (MAPS), to be explored and developed for the current application.

2.1 TARGET/BACKGROUND PARTITION - THE REARC TECHNIQUE

Redundant Area Coding (REARC) partitions the original (2048)² scene into an 8 x 8 matrix of (256)² pixel subframes. A simple 64-element binary matrix is then used to distinguish subframes which are in the "target" (ones) from those in the "background" (zeros). The background is used only to provide context for the target regions and is thus allowed to be presented at significantly degraded fidelity relative to the original image.

The compression ratios in the target (C_T) and background (C_B) required to meet a given overall compression ratio (C_O) are related as shown in Equation (2.1):

$$\frac{N_T}{C_T} + \frac{64 - N_T}{C_B} = \frac{64}{C_O} \quad (2.1)$$

Here, N_T denotes the number of REARC subframes (out of 64) which have been designated as target regions. This relation is plotted in Figure 2-1 with C_0 fixed at 30:1 and N_T parametric at values of 4, 6, and 8 target subframes.

Several observations may be drawn from these plots. First, it is seen that the required background compression rises asymptotically as the target compression approaches a specific minimum value determined by the number of target subframes, N_T . In essence, all of the bits are used up in coding the target at the minimum C_T and none remain for the background region. Second, these curves represent relatively small target fractions - only one-eighth of the scene at $N_T = 8$ and one-sixteenth at $N_T = 4$. Even so, significant minimum target compressions are required if $C_0 = 30$ is to be obtained. Third, any practical representation of the background regions will probably require compressions of no more than about 60:1. Thus, target compressions in the range of 4:1 up to 10:1 will be needed.

Since non-destructive coding is limited by the entropy of typical imagery to values in the range of two-to-four bits per pixel or compression ratios in the 1.5:1 to 3:1 range, the data in Figure 2-1 imply the need for destructive coding in both target and background regions.

2.2 ALGORITHM CLASSIFICATION

Available compression algorithms divide quite naturally into three general classes - entropy, spatial, and transform encoding.^{2,3,4}

2.2.1 Entropy Encoding

Entropy encoding seeks to reduce the redundancy in an image description by redefining the "alphabet" describing the image intensities on the basis of the statistical occurrence of the various "symbols". Short symbols are used for frequently occurring values and larger bit strings for those which show up more rarely. Huffman coding is the outstanding example of this approach. Such techniques rarely reach 2:1 compression ratios and more typically give a few tens-of-percent improvement.




Figure 2-1. Target/Background Compression Requirements

Note that in theory it is possible to apply entropy encoding to the "symbol" stream which results from some other compression technique (spatial or transform) and thus achieve some further reduction. However, most high compression techniques inherently remove redundancies associated with frequently observed patterns in the original imagery and thus tend to "flatten" the statistical distribution. In this case, the opportunities for further gains based on a revision of the symbol alphabet are very slight. Hence, entropy encoding techniques are too compression-limited to be of much assistance in the current application.

2.2.2 Spatial Encoding

The spatial approach to encoding includes several classical techniques. This class further subdivides into "integrating" and "differential" subclasses, classical run-length encoding performed on either one or two-dimensional sequences is probably the best known representative of the integrating technique subclass. Here compression is achieved, as the name implies, by giving the intensity value and a code describing the length of the "run" of contiguous pixels which share that intensity. This technique exploits the strong correlation between adjacent pixels which is exhibited over much of a typical image. Effectively, the encoding integrates the description of several pixels.

The differential encoding subclass of the spatial techniques is best represented by the classical DPCM technique and its many variants employing both fixed and adaptive coding strategies. Again, the strong element-to-element correlation is exploited but in this case by describing the "differential" step between elements. Since the correlation renders these steps typically small, compression is achieved because fewer bits are required to describe the differences than were needed to describe the full intensities.

These classical spatial techniques work well (in the adaptive modes) down to about two bits per pixel or compressions of 3:1. They can be further pushed to the neighborhood of one bit per pixel or compression ratios of 6:1 with modestly good fidelity but generally degrade rapidly beyond this level.

Since the goal is an overall compression to 0.2 bits per pixel, the classical spatial techniques have marginal utility in the present application.

A new two-dimensional integrating spatial encoding technique denoted as Micro-Adaptive Picture Sequencing (MAPS) is found to provide good fidelity performance out to the required compression levels. MAPS is described in detail in section 3; key characteristics which led to its ultimate selection for this application are summarized here:

- MAPS retains suitable fidelity over required compression ranges for both target and background regions;
- MAPS is compatible with the REARC partition and represents a single technique with only control parameter changes between target and background;
- MAPS has a very fast and compact all-integer recursive implementation;
- MAPS has inherent detection (and potential correction) capability to prevent propagation of the effects of transmission line bit errors.

2.2.3 Transform Encoding

Linear transforms represent the final broad class of available coding techniques and a wealth of these have been reported in the literature. Among the transforms applied and described are the Walsh-Hadamard, Haar, Slant, Discrete Linear Basis, Discrete Fourier, Discrete Cosine, and Karhunen-Loève.⁵ The Walsh-Hadamard Transform (WHT) has the simplest digital implementation. The Karhunen-Loève Transform (KLT) lies at the other end of the complexity spectrum. In essence, the KLT expands the image in terms of itself and thus gives the best fidelity performance at a given compression level. The remaining techniques fall between the WHT and KLT in fidelity and complexity. The image compression "community" seems to have converged upon the Discrete Cosine Transform (DCT) as yielding the best compromise. Its fidelity performance approaches that of the KLT but its implementation complexity is much more modest.

Unless a rather elaborate image adaptation strategy is added to the already significant complexity of the basic transform implementation, this whole class of techniques exhibits reasonably good performance down to a little under one bit per pixel and rapid degradation when attempts are made to press further. Qualitatively it appears that this degradation arises in a rather subtle manner. It is characteristic that a transform technique builds up the details in an image by superposition of contributions from several of its basis vectors. The strong details in the image result in large coefficient amplitudes for the corresponding basis vectors and the compression strategy seeks to preserve these "high energy" terms while retaining only a minimal representation of the small coefficients. This approach does indeed retain the strong details but it also "aliases" them into other parts of the image block over which the transform was applied. In the complete transform description with all terms accurately represented, these aliases are suppressed by destructive interference from the myriad low energy terms. Note that the strong detail is also sharpened by terms of this type. Rapid degradation in transform coding then results at that point where it becomes necessary to completely eliminate information from the several low energy terms in order to achieve further compression.


The problem just described is most severe with a fixed coefficient selection scheme and can be somewhat alleviated by resorting to an image-content-adaptive scheme. However, just the operation of taking the transform (or its inverse) typically requires on the average of one hundred instructions per pixel. If the overhead associated with the adaptation is added to this, the transform technique appears to require about an order-of-magnitude more operations than does MAPS!

2.3 CONCLUSION

MAPS and the adaptive transform techniques (specifically the adaptive DCT) are the only approaches uncovered which show the potential for meeting the compression goals. Moreover, the MAPS technique exhibits several alternatives within its basic conceptual framework. These options coupled with the overwhelming disparity in implementation complexity between MAPS

and the transform approach led to the decision to invest the remaining effort in exploring and developing some of the intra-MAPS trade-offs. The classification net leading to this selection is summarized in Figure 2-2.

The next section is devoted to a self-contained description of MAPS and the specific intra-MAPS trade-offs are presented in Section 4.



ALGORITHM SURVEY

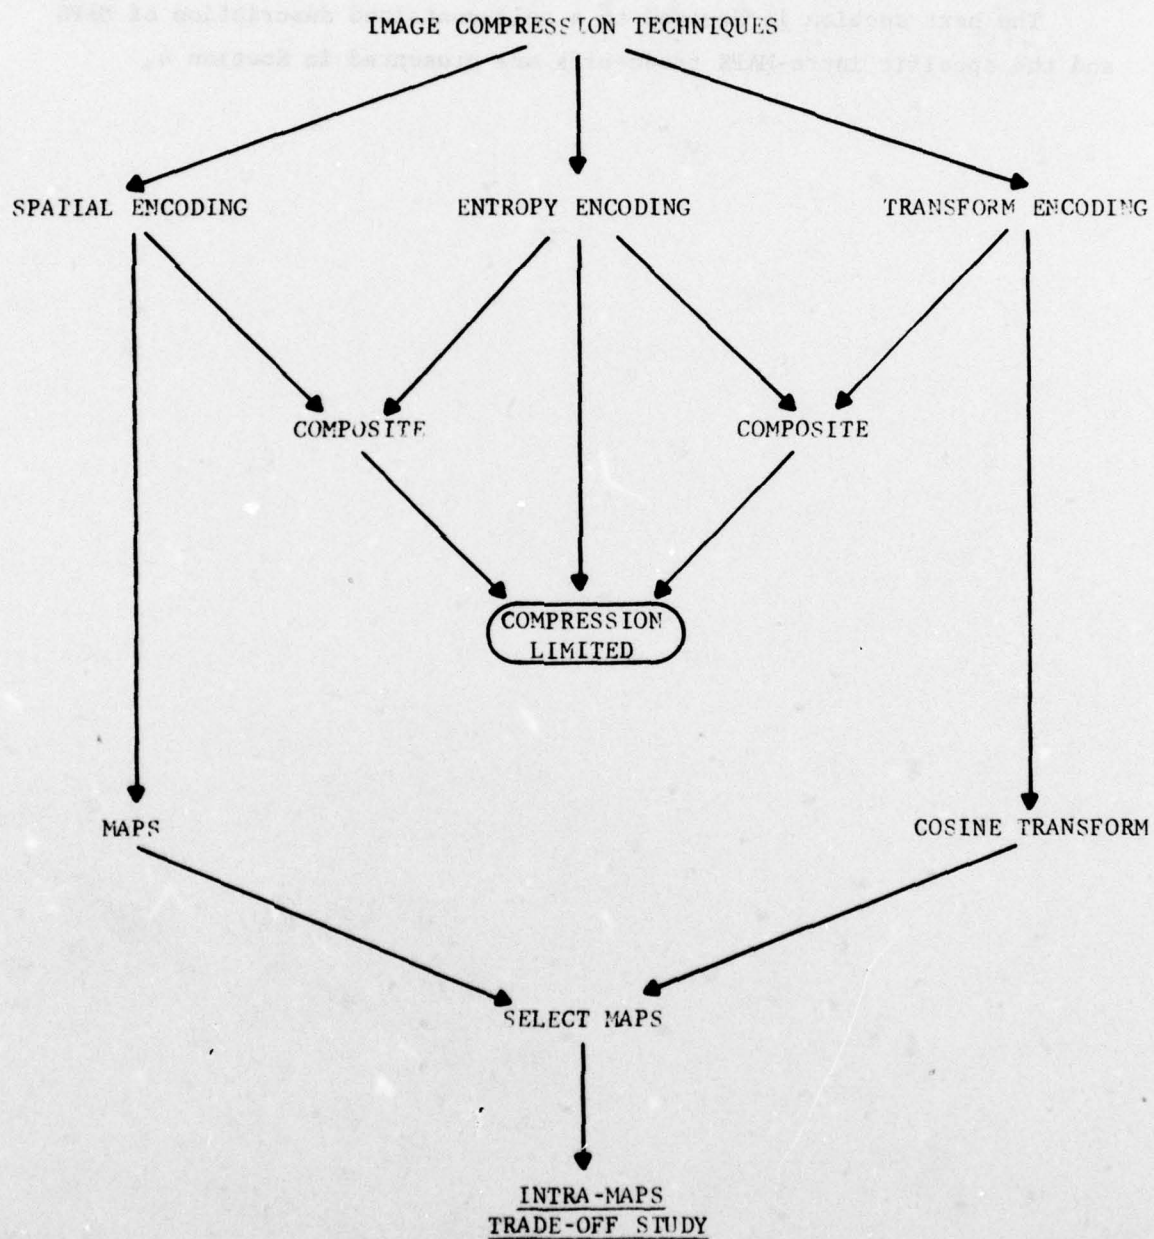


Figure 2-2. Compression Algorithm Classification Net

SECTION THREE

TWO-DIMENSIONAL IMAGE CODING BY MICRO-ADAPTIVE PICTURE SEQUENCING (MAPS)

Micro-Adaptive Picture Sequencing (MAPS), a computationally-efficient contrast-adaptive variable-resolution digital image coding technique, is described in this section. Both compression and decompression involve only integer operations with no multiplies or explicit divides. The compression step requires less than 20 operations per pixel and the decompression step even fewer. MAPS is based on the combination of a simple vision heuristic and a highly nonlinear spatial encoding. The heuristic asserts that the fine detail in an image is noticed primarily when it is sharply defined in contrast while larger more diffuse features are perceived at much lower contrasts. The coding scheme then exploits the spatial redundancy implied by this heuristic to maintain high resolution where sharp definition exists and to reduce resolution elsewhere.

In the 'standard' digital encoding of an image, the scene is partitioned into a grid of equi-sized pixels and an intensity value is given sequentially for each element. Thus, only the intensity data is explicit. The resolution and position data are implicit - the resolution because the pixel size is constant and the position because it has a trivial mapping with the sequential location of the corresponding intensity value in the data stream. The resolution is chosen in order to preserve the smallest detail deemed significant in the image. However, this resolution is often required only for small regions of the image and much coarser resolution would suffice in other areas. From this standpoint, the image encoding problem becomes one of finding a compact method for describing the varying resolution and an automatic method for selecting the appropriate resolution locally.

Micro-Adaptive Picture Sequencing (MAPS) departs from 'standard' matrix encoding of the image by making the resolution explicit (See Figure 3-1a). Indeed, this is done at a micro level and element size can vary from 'pixel' to 'pixel'. Furthermore, this process is organized in such a way that the

'STANDARD' ENCODING:

MAPS ENCODING:

INTENSITY RESOLUTION POSITION

EXPLICIT	IMPLICIT	IMPLICIT
EXPLICIT	EXPLICIT	IMPLICIT

(a) image description




Figure 3-1. The Bases of MAPS

MAPS element positions remain implicit. Thus, compression is achieved by replacing a large amount of (redundant) intensity data with a smaller amount of explicit resolution information.

Solution of the second part of the problem - automatic control of local resolution in MAPS - is based on a simple vision heuristic: 'When an image is viewed as a whole, fine detail is noticed only when it exhibits sharp contrast'. This concept is extended and presented schematically in Figure 3-1b. Detail which falls below the threshold locus in this contrast-size space is assumed to be ignorable. Thus, portions of the image which fall in this region can be presented at lower resolution. In practice, resolution is decreased until it reaches its last value prior to crossing the locus.

The remainder of this section describes how these simple basic concepts are coupled to provide image compression and decompression. In addition, the compatibility of MAPS with a system of interactive macro-fidelity control is discussed. Finally, several examples illustrating MAPS performance and its various facets are presented.

3.1 MICRO-ADAPTIVE PICTURE SEQUENCING

Computationally, MAPS has an extensive decision structure; it tends to be logically complex but arithmetically very simple. Thus, it will be described in terms of the procedures required.

3.1.1 Compression

The details of MAPS compression are to be understood in terms of the answers to three questions:

- How is element position described implicitly?
- How is the local resolution determined?
- How is the local resolution encoded?


These three are dealt with in the following three subsections.

3.1.1.1 Sequence Convention

The element position can be inferred from the past history of the data stream provided an unambiguous order convention is present. MAPS is applied to successive independent image subframes of size $2^M \times 2^M$. Each subframe is partitioned by successive two-dimensional binary subdivisions until the original resolution (1×1 or $2^0 \times 2^0$) is reached. The 'level' L of the elements in a particular subdivision is just the exponent which gives the element edge 2^L . Thus, there are M subdivisions and $M + 1$ levels running over $0 \leq L \leq M$. For any local subdivision of an element at level $(L + 1)$ into four elements at level L , the local position order is always that given in Figure 3-2a for any value of L . The nested elements for $M = 3$ with their local position designations (including 'level' as a subscript) are shown in Figure 3-2b. The sequence convention starts at the lower left corner (local position zero for all levels) and simply increments local position in the level hierarchy from 0 to $(M - 1)$ until the upper right hand corner is reached.

The path through the subframe which is implied by this sequence is illustrated in Figure 3-2c. Note that the local position at each level is a two-bit quantity in which the lower bit indicates the local 'x' position and the upper bit the 'y' value. A little thought should suffice to observe that these two bits are just the L th bits, respectively, of the full intrasubframe 'x' and 'y' addresses. Thus, the indexing of the image array in this zig-zag order is essentially automatic.


Note that the order is such that all subblocks of a block at a given level and local position are traversed before any portion of a block at the same level but a higher local position is encountered. This principal of 'block exhaustion' is the key order constraint which allows MAPS to work as a single-pass compression procedure and at the same time retain the position implicitly. This particular order is denoted as the 'block' sequence.



(a) local position numbering
(all levels $0 \leq L \leq M$)

2_0	3_0	2_0	3_0	2_0	3_0	2_0	3_0
-2_1	-3_1	-2_1	-3_1	-2_1	-3_1	-2_1	-3_1
0_0	1_0	0_0	1_0	0_0	1_0	0_0	1_0
2_0	3_0	2_0	3_0	2_0	3_0	2_0	3_0
-0_1	-1_1	-0_1	-1_1	-0_1	-1_1	-0_1	-1_1
0_0	1_0	0_0	1_0	0_0	1_0	0_0	1_0
2_0	3_0	2_0	3_0	2_0	3_0	2_0	3_0
-2_1	-3_1	-2_1	-3_1	-2_1	-3_1	-2_1	-3_1
0_0	1_0	0_0	1_0	0_0	1_0	0_0	1_0
2_0	3_0	2_0	3_0	2_0	3_0	2_0	3_0
-0_1	-1_1	-0_1	-1_1	-0_1	-1_1	-0_1	-1_1
0_0	1_0	0_0	1_0	0_0	1_0	0_0	1_0

(b) level nesting, $M = 3$



(c) 'block' sequence path
(original image)

Figure 3-2. Original Image Subframe Partition


The MAPS description of a subframe can consist of any complete, non-overlapping set of the naturally-arising blocks at various levels. Figure 3-3a shows a sample MAPS spatial partition of a subframe with $M = 3$. Figure 3-3b, on the other hand, shows a configuration which is not allowed because the position integrity is lost. In effect, the pattern in Figure 3-3b violates the block exhaustion constraint.

In essence, block exhaustion imposes a dynamic local maximum block size constraint. Once a block at level L has been left in a subdivided state, no blocks larger than L can be formed until a block at local position 3_L has been completed. Furthermore, completed MAPS subblocks at level L (and greater) which were being held pending the possibility of combining them into even larger blocks can be disposed of as final MAPS elements as soon as the dynamic local maximum level drops to L or smaller. Thus, only a very small amount of intermediate storage (enough for up to three pending blocks at each level from $L = 1$ to $L = M - 1$) is required.


Once the MAPS partition for a full subframe has been completed, the order of the data is just that implied by the block sequence convention. The path implied by the partition shown in Figure 3-3a is presented in Figure 3-3c. This illustrates the completely implicit nature of the MAPS element positions, the answer to the first question.

3.1.1.2 Contrast-Adaptive Resolution Control


The local contrast of the image is evaluated only among four complete subblocks covering the natural block at the next higher level. Denote the intensities of these four subblocks as I_0, I_1, I_2, I_3 where the subscript characterizes the local position. The intensities are sorted into increasing order and differences are then taken between the successive values and between the highest and lowest intensity. These are designated as the 'lower step', 'middle step', 'upper step', and 'extreme' contrasts respectively (See Figure 3-4a).



(c) Block sequence path



(b) An excluded configuration



(a) A possible MAPS configuration

Figure 3-3. Sample MAPS Subframe Partition.

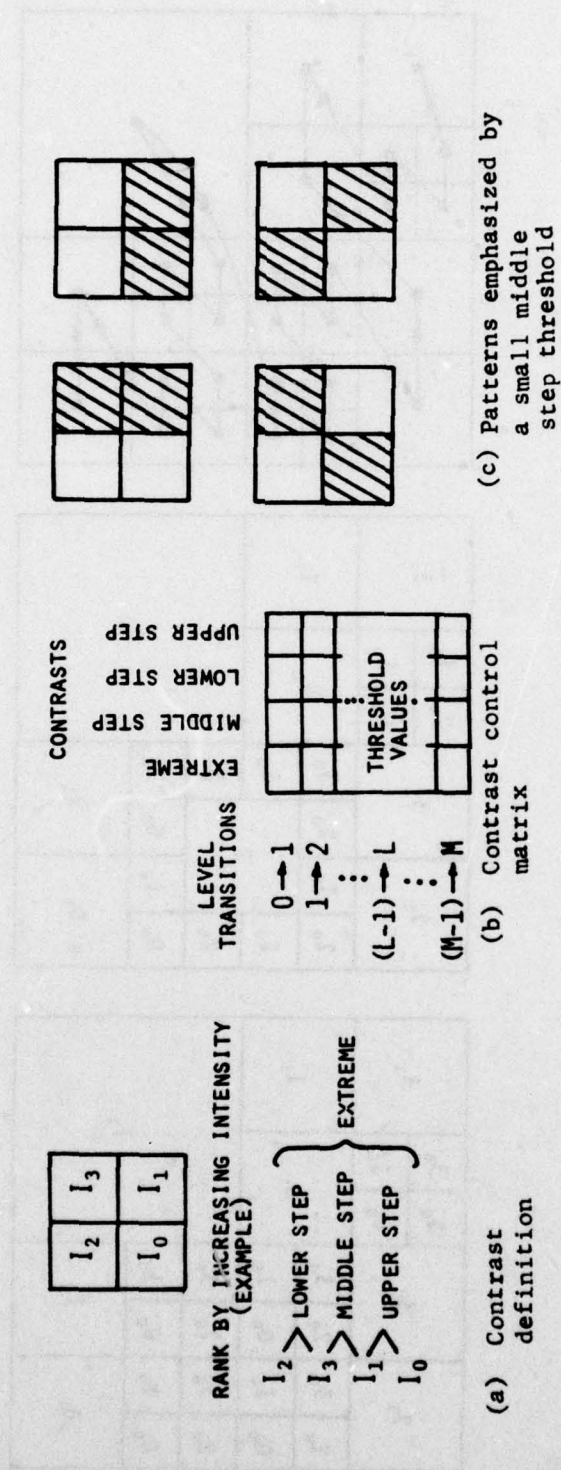


Figure 3-4. Micro-Fidelity Control

If the level of the four subblocks is $L-1$ then each contrast is tested against a separate threshold corresponding to (1) its contrast type and (2) the transition from level $L-1$ to level L . Thus, the contrast thresholds form a contrast control matrix as shown in Figure 3-4b.

When none of the four thresholds are violated, the four subelements are combined into a single new element at the next higher level. This element is either (1) retained as a pending subblock for the level above it or (2) disposed of as a final MAPS element if its level coincides with the current local maximum level. The combination step is a simple averaging of the four subblock intensities. This has the virtue that it minimizes the mean square error between the four source elements and the composite.

In the alternate case, when one or more of the thresholds is violated, the four subelements become final MAPS elements (at level $L-1$), the local maximum level is reset to L , and any pending elements are also finalized.

Each column of the contrast control matrix is typically chosen to give a discrete approximation to a threshold locus similar to the one depicted in Figure 3-1b. The entries in each row are typically chosen so that the middle step threshold is the smallest, and the lower and upper step thresholds are smaller than the extreme. The latter, in turn is usually smaller than the sum of the three step contrast thresholds.

The smaller value for the middle step threshold is predicated on a desire to preserve patterns of the form shown in Figure 3-4c. These might be part of a faint but extended linear feature which is often of interest.

3.1.1.3 Explicit Resolution Codes

The block size or resolution code of each MAPS element is uniquely specified by its 'level'. Thus, a two-bit code can cover subframes up through $M=3$ and a three-bit code is sufficient through $M=7$. Note, however, that the level is just \log_2 (block edge) or \log_4 (block pixel count). Two bits will cover a block size dynamic range of 64 to 1 and three bits could cover up to 16384 to 1! This \log_4 coding is very compact.

The selection of the subframe or maximum block size ($2^M \times 2^M$) is a compromise between achievable compression and the memory required for temporary storage of the original image. Empirical experience with a fairly broad range of images has shown that little is gained beyond $M = 4$ (or 16×16). The few larger blocks which could show up in a typical image make a negligible difference in the compression (a few percent at most) but cost enormously in memory which doubles for each increment in M .

If compressions are not to be pressed below about 0.5 bits per pixel, then a choice of $M = 3$ is to be preferred. In this curious situation both compression and mean square error improve by reducing the subframe size to 8×8 . The compression improves because the resolution code overhead drops from three bits to two bits per MAPS element. The fidelity can improve because blocks of even slightly different intensity at level 3 are no longer averaged to level 4.

For either choice of subframe size, the MAPS element or 'pixel' is thus a compound code consisting of a two or three-bit resolution code and an intensity code of as many bits as desired (usually six or eight).

3.1.1.4 Implementation

From the description above, it is seen that the MAPS procedure is all integer. There are no multiplies required. Moreover, since the divisor needed to form the composite intensity is always four, this 'divide' operation can be implemented by adding two to the four-element intensity sum (to round it) and then performing a simple two-bit right shift. Thus, the 'instruction set' for MAPS is very small - mainly adds, shifts, and conditionals. With the use of the adder, the conditionals can all be trivially changed to a sign-bit jump.

These limited and very rapid instruction types result in very fast compression with MAPS. Detailed empirical operations counts have been performed over a wide range of imagery types and compression ratios. Here, 'operations' include such steps as local register fetches and index updates

as well as the more usual steps such as adds and image memory references. A single 'instruction' therefore covers two to four 'operations'. The observed counts are consistently under 20 operations per original pixel. The counts depend only slightly on the compression ratio achieved even though the procedure involves a variable pattern of recursion.

The entire compression procedure is summarized in the MAPS subframe block compression functional flow which is displayed in Figure 3-5. The process is highly recursive so the operational code is itself very compact. The recursion for levels greater than $L = 0$ is shown as in the inner loop in Figure 3-5.

3.1.2 Decompression

Two decompression modes have been explored for MAPS image descriptions.

3.1.2.1 Block Mode

The block mode of decompression is trivial and extremely fast. The image address is merely incremented according to the block sequence convention and each original matrix pixel is set equal to the intensity of the current MAPS element. The number of pixels in the current element is found from the resolution code through a simple table lookup.

MAPS overall speed and the further enhancement provided by the compression/decompression speed asymmetry may be effectively exploited in a data base environment. In such applications, the image is often compressed once but then accessed (decompressed) many times. If the image is to be used only for display however, the block mode has the disadvantage that some elements are large enough to make the boundaries noticeable through the Mach effect. The resulting 'blockiness' is similar to that observed with many compression techniques which process elements in independent subframes. The technique described next helps alleviate this problem.

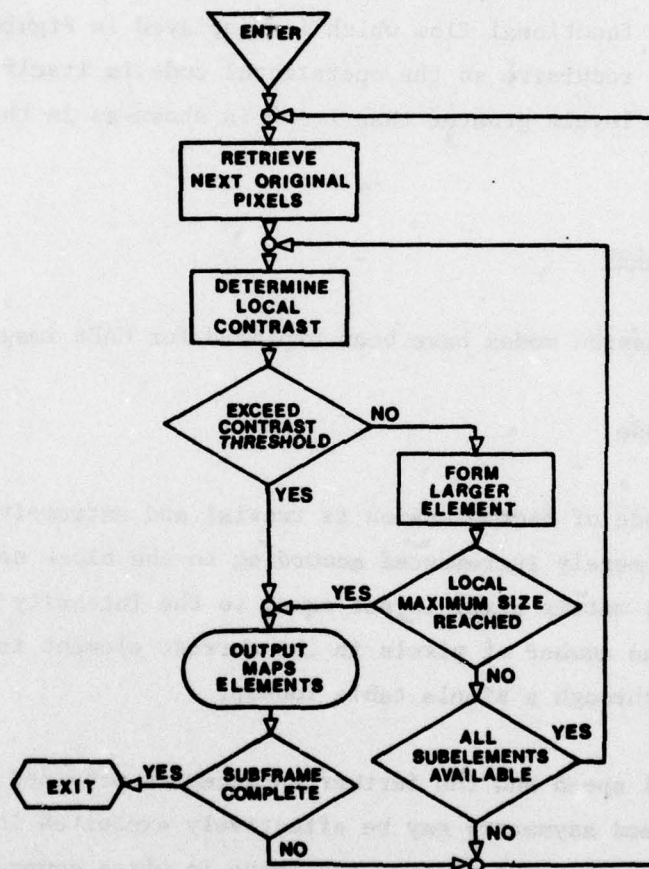


Figure 3-5. MAPS Subframe Block Compression - Functional Flow

3.1.2.2 Adaptive Mode

The size of each MAPS element and those which surround it gives an indication of the detail present in that local area. This information can be used directly to guide an adaptive convolution process over the MAPS image during decompression. When the pixels within the region covered by a given MAPS element are to be set, the convolution window is chosen to be one pixel narrower than the MAPS element. This makes the window odd and thus symmetric about the target pixel. The choice of window size is based on the idea that the MAPS element size is an estimate for the local correlation length.

Next, the size is determined for each of the MAPS elements covering the corner pixel in each of the twelve surrounding regions shown in Figure 3-6. The MAPS information from each of these surrounding blocks is used in the convolution only if the condition $L_{\text{surround}} \geq L_{\text{central}} - 1$ is met. There is a two-fold motivation for this. First, an element in the 'surround' which is more than one level smaller than the central element is an indication of higher activity along that edge. Second, if the 'surround' element is no more than one level smaller than the central block, it is guaranteed that the MAPS intensity is uniform within that block of the surround. This, in turn, means that there are, at most, thirteen independent inputs over the entire central-element convolution. An efficient implementation can be established on this basis.

The convolution process removes the 'blockiness' effectively. In large, almost-flat regions, however, the images will still exhibit 'contouring.' Addition of a small random dither during the convolution effectively obliterates these remaining artifacts.

Even with the dither, the adaptive decompression not only improves the 'looks' of the image but empirically it has consistently produced a smaller mean square error than that given by the block mode!

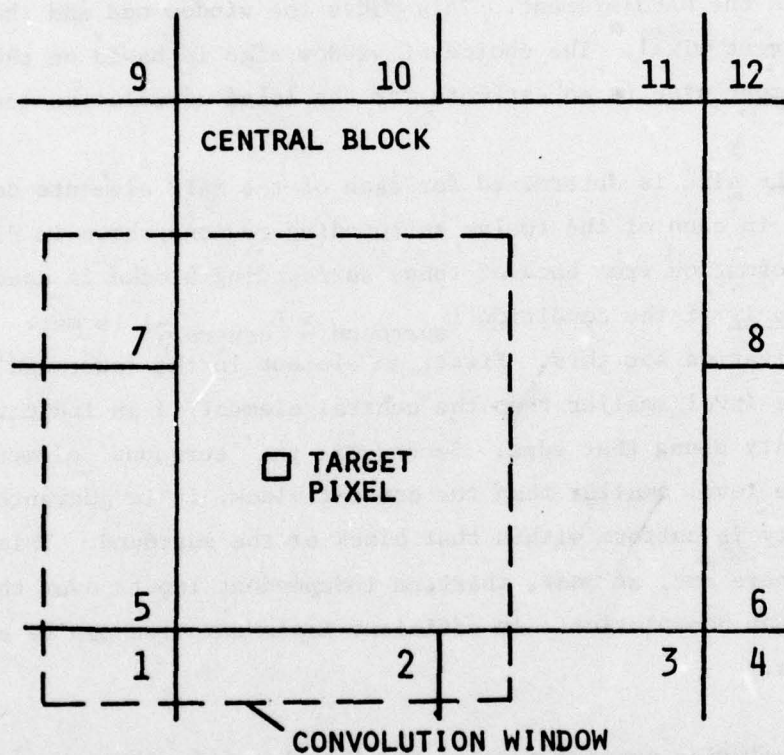


Figure 3-6. Adaptive Decompression Geometry

3.1.3 Interactive Macro-Fidelity Control

MAPS operates independently on each successive $2^M \times 2^M$ subframe within the image. As a consequence, it is directly compatible with systems which employ a macro-fidelity criterion which varies over the scene. The REDundant AREA Coding scheme of the Experimental Image Compression System (EICS)¹ is an example of such selectable fidelity control. In EICS, the overall scene is partitioned into an 8×8 grid of subimages. Then some subset of the 64 subimages is 'spotlighted' for preservation at high fidelity while the remaining 'background' is needed only for context and is allowed much more degradation.

With MAPS, the spotlight and background regions are distinguished simply by using low contrast control thresholds in the high resolution area and much higher thresholds in the background. Control is effected merely by altering the selection of one of two contrast control matrices as directed by the REARC partition. Here it is understood that each subimage (typically 256×256) is covered exactly by a set of several MAPS subframes (typically 16×16 each).

The range of control which may be selected through specification of the contrast control matrix is very broad. It starts with perfect fidelity when all contrast thresholds are set to zero. In this case, the natural redundancy in the image will still allow some composite blocks to be formed. However, the overhead associated with the explicit resolution codes will usually keep the compression ratio near unity with either a slight compression or expansion depending on the particular image. At the other end of the control range, a forced image reduction to any given MAPS level can be achieved by setting all contrast thresholds up to that level equal to the maximum intensity plus one. In this case, it is impossible for the thresholds to be violated. Forced reduction in the background region is an effective strategy when very high overall compressions must be reached.

Note that the MAPS information is entirely local and is complete when the image is covered with blocks of explicitly given resolution and intensity. No additional information describing the controls used in arriving at the MAPS data stream is required for decompression. In this sense, MAPS decompression is 'transparent' to any variations in control which were imposed during the compression phase.

The transparency just described forms the basis for a much more elaborate strategy of interactive macrofidelity control. The user could now 'spotlight' in regions of nearly arbitrary boundary made up of an integral number of MAPS subframes. In fact, the partition could extend to regions within regions giving a kind of foveal view in the spotlight. A different contrast matrix may then be used for each region if desired. Any feature in the scene could thus be kept to any required fidelity. The strength of the approach then lies in the fact that no overhead information need be carried to describe this control 'scaffolding' in the final MAPS description.

3.2 MAPS EXAMPLES

MAPS performance is illustrated by the several images in this section. The IEEE Facsimile Test Chart and the Airfield scene were each digitized in a 2048 x 2048 matrix and the power plant scene was given by a 1920 x 1856 matrix. A maximum MAPS element level of $M = 4$ (16 x 16 subframes) was used throughout. Each resultant image is characterized in its caption with the relevant decompression mode, macro fidelity control mode, compression in bits per original pixel, and/or mean square error.

3.2.1 IEEE Facsimile Test Chart

The IEEE Facsimile Test Chart results are presented in Figure 3-7. This example illustrates that MAPS is applicable to diverse imagery types and that annotation on an image can be included naturally in the MAPS data stream.

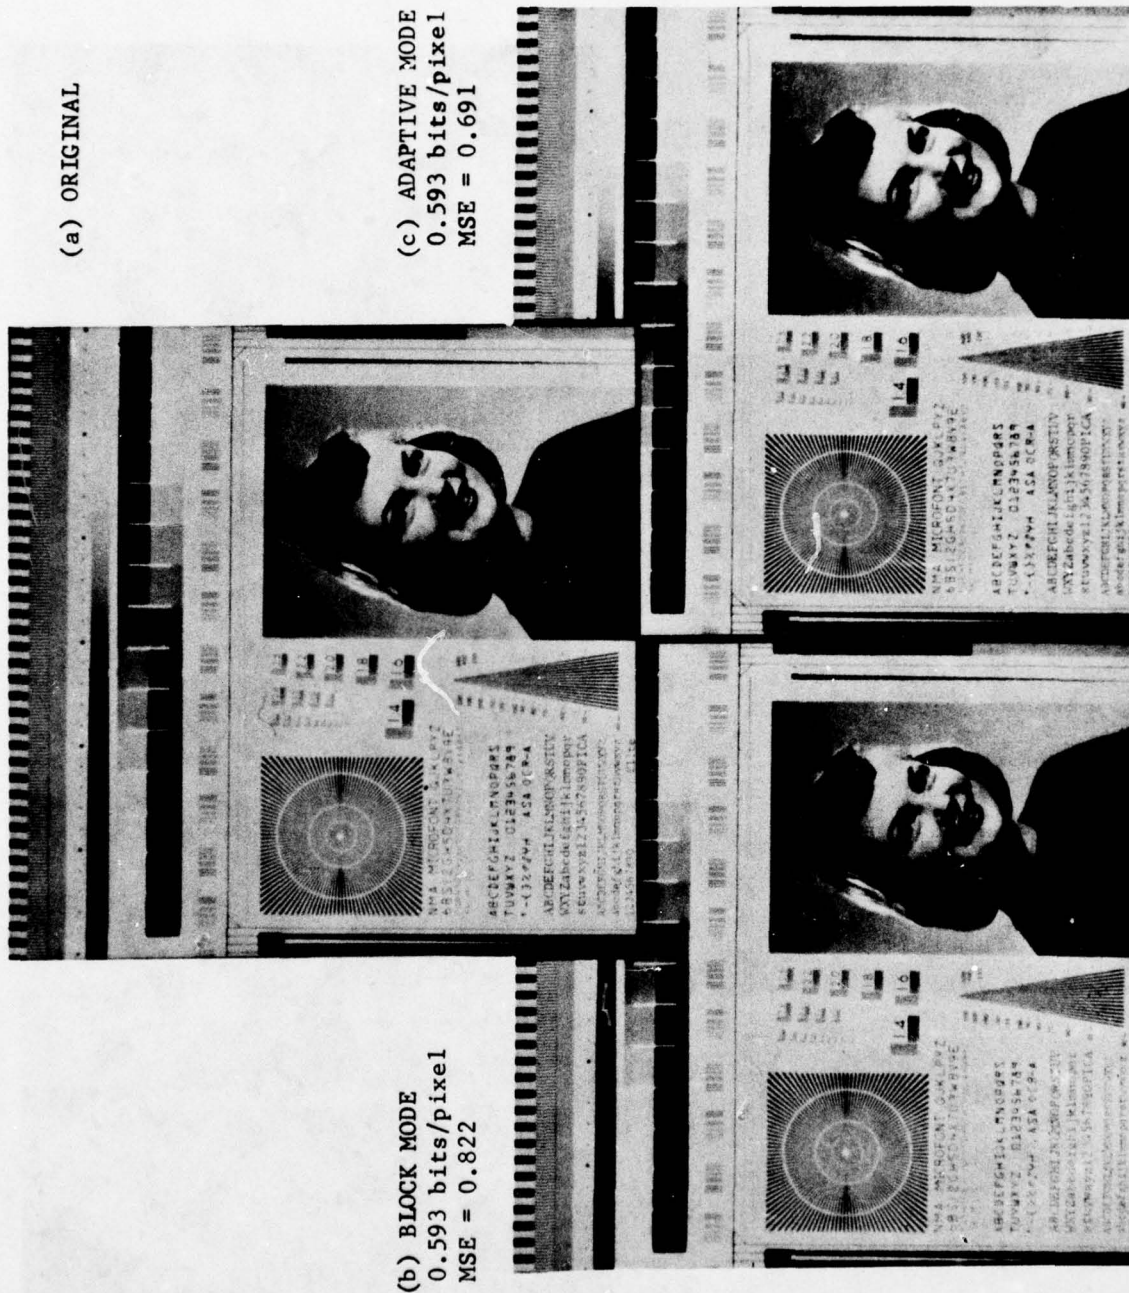


Figure 3-7. IEEE Facsimile Test Chart





Figure 3-8. Airfield Scene



(a) POWER PLANT SCENE
(Spotlighted Region)
Six-bit ORIGINAL

(b) 0.167 Bits/Pixel OVERALL
0.762 Bits/Pixel SPOTLIGHT
0.119 Bits/Pixel BACKGROUND

Figure 3-9. Power Plant Scene

The artifacts associated with the block decompression mode (Figure 3-7b) are particularly noticeable in the girl's face. The adaptive decompression mode (Figure 3-7c) eliminates this effect and also reduces the mean square error as noted earlier.

3.2.2 Airfield Scene

MAPS results for the Airfield scene are presented in Figure 3-8. The compression to 0.455 bits per pixel (Figure 3-8b) is included to illustrate the application of a fixed control matrix over an entire tonal image.

With a spotlight/background distinction, the overall compression is extended to 0.198 bits per pixel (Figure 3-8c). Definition in the background region is clearly still sufficient for context purposes.

3.2.3 Power Plant Scene

The final example is given in Figure 3-9 and shows another spotlight/background application - this time with an overall compression to 0.167 bits per pixel.

In all of these examples, the number of operations per pixel in the compression steps is less than 20. Thus, MAPS remains computationally efficient to very high compressions. This efficiency should allow MAPS to be effectively implemented in a variety of real-time interactive-control applications for image storage and transmission.

SECTION FOUR

INTRA-MAPS TRADE-OFF STUDIES

The basic MAPS concepts provide the framework for a robust set of options which impact the implementation speed and cost, compression/fidelity relationship, and error detection/correction potential. Three of these selections are studied in detail:

- Contrast Control Mode
- Highest Resolution Code
- Maximum Block Size.

4.1 COMPOSITE TEST IMAGE

The full test imagery set for this effort consists of ten 2048 x 2048 scenes covering a diverse range of image types. Within these full frame scenes, a total of forty-four 256 x 256 subframes were pre-designated as REARC target regions. This subset was augmented by four additional subframes to compose an 8 x 6 matrix of 'representative' imagery as exhibited in Figure 4-1. These target subframes actually tend to be somewhat 'busier' than the typical subframe so the composite presents a more-difficult-than-average compression challenge.

A range of MAPS contrast control matrices yields the empirical fidelity versus compression relationship shown in Figure 4-2 for the overall composite. The mean square error defined as

$$MSE = 100 \frac{\sum (I_{MAPS} - I_{ORIGINAL})^2}{\sum (I_{ORIGINAL})^2} \quad (4.1)$$

is used here as the fidelity metric. The quantity I_{MAPS} represents the individual pixel intensity in the MAPS block-decompressed image, $I_{ORIGINAL}$




Figure 4-1. Composite Test Image - Original




Figure 4-2. Fidelity vs Compression - 48 Subframe Composite Image

is the original image pixel intensity, and summation extends over the full composite. Note that the MSE reaches 0.5% at a compression level of less than one bit per pixel! Given the diversity and difficulty of this composite, the overall MAPS performance holds up very well.

The composite image compressions form the empirical base for the MAPS trade-off selections described in the next three subsections.

4.2 SELECTION OF MAPS CONTRAST CONTROL MODE

As described in Section 3.1.1.2 and illustrated in Figure 3-4, each MAPS level transition is controlled by tests on four contrast values denoted as 'extreme', 'lower step', 'middle step' and 'upper step'; this four-contrast case will be referred to as the 'step' contrast control mode. Alternatively, control can be effected on the 'extreme' contrast alone and the contrast control matrix shrinks to a control vector; this will be denoted as the 'extreme' mode.

The trade-off here is between fewer operations and simpler contrast threshold specification for the extreme mode versus greater flexibility in fidelity specification for the step mode. The extreme mode requires one less conditional in determining the contrast and up to three fewer conditionals in testing the contrasts against the corresponding transition thresholds. Each of these extra conditionals requires several related operations such as differencing intensities or retrieving the test threshold from a local register file. However, the operations 'cost' of these several steps is amortized among four elements at the current MAPS level. Thus, the impact in terms of operations per original pixel is not as severe as it would at first appear.

The qualitative fidelity advantage of the step mode resides in its ability to preserve horizontal and vertical edges or diagonal lines by applying a tight threshold to the 'middle step' contrast (see Figure 3-4c). This allows relaxation of the 'extreme' threshold in the step mode which in turn permits smoothing over a wider dynamic range of isolated local noise fluctuations

than was possible in the extreme-only mode.

The competing effects of this trade-off are illustrated in Figures 4-3 (extreme contrast) and 4-4 (step contrast). Note that the 'a' portion of each figure provides a map of the compression/MSE/operations count for each individual subframe of the composite. The overall performance values and the control matrix are summarized at the top of this same map. The overall compression ratios are seen to be essentially identical for the two cases (just under 5:1) while the extreme mode exhibits somewhat fewer operations per pixel (13.4 versus 17.9) as expected.

It is also apparent that the extreme mode has slightly better mean square error performance. On reflection, this is not too surprising since the extreme contrast threshold in the step mode is relaxed relative to that in the extreme (-only) mode and a wider dynamic range of noise is smoothed in the step contrast case. The improvement in fidelity for the step mode comes in the structured portion of the image and the MSE does not reflect this advantage. The fidelity improvement is perhaps best illustrated in the narrow roads surrounding the air strip in subframe 4. Visual inspection of this region shows a clearcut advantage for the step mode. Note that for this particular subframe, the compression in the step mode is also higher than that in the extreme mode.

Operations counts for the step mode increase only by about one third relative to the extreme mode. Further, the operations count is still very low in an absolute sense for any technique operating at high compression ratios. Since a clearly discernible fidelity improvement and greater fidelity specification flexibility are available in the step mode, the small penalty in increased operations count is accepted.

4.3 SELECTION OF MAPS HIGHEST RESOLUTION CODE

Each MAPS 'element' provides a compound description of the local resolution MAPS 'level' or \log_4 (pixel count per block) and intensity. The resolution code portion

EXTREME CONTRAST - LMAX=3

C= 4.98, MSE= .331, CPS=13.4, CONTROL 8 8 8 8
5 5 5 5
3 3 3 3
0 0 0 0
0 0 0 0
0 0 0 0

SUBFRAME	I 8	I 16	I 24	I 32	I 40	I 48	I
COMPRESSION	I 2.81	I 9.71	I 5.24	I 8.73	I 5.31	I 11.29	I
MSE	I .281	I 1.234	I .194	I 1.249	I .846	I .211	I
OPS/PIXEL	I 13.18	I 13.54	I 13.42	I 13.54	I 13.40	I 13.59	I
SUBFRAME	I 7	I 15	I 23	I 31	I 39	I 47	I
COMPRESSION	I 3.80	I 7.80	I 5.35	I 7.25	I 11.50	I 4.30	I
MSE	I .792	I .778	I .111	I .719	I .585	I .422	I
OPS/PIXEL	I 13.29	I 13.51	I 13.44	I 13.50	I 13.55	I 13.31	I
SUBFRAME	I 6	I 14	I 22	I 30	I 38	I 46	I
COMPRESSION	I 6.08	I 4.32	I 8.19	I 6.06	I 14.20	I 3.73	I
MSE	I .733	I .655	I .142	I .437	I .219	I .467	I
OPS/PIXEL	I 13.42	I 13.30	I 13.51	I 13.45	I 13.57	I 13.23	I
SUBFRAME	I 5	I 13	I 21	I 29	I 37	I 45	I
COMPRESSION	I 5.90	I 3.82	I 7.39	I 4.53	I 12.07	I 1.34	I
MSE	I .909	I .951	I .116	I .181	I .149	I .236	I
OPS/PIXEL	I 13.41	I 13.26	I 13.48	I 13.34	I 13.56	I 12.76	I
SUBFRAME	I 4	I 12	I 20	I 28	I 36	I 44	I
COMPRESSION	I 5.73	I 3.74	I 5.11	I 5.01	I 23.17	I 7.21	I
MSE	I .861	I .607	I .130	I .381	I .108	I .258	I
OPS/PIXEL	I 13.40	I 13.21	I 13.42	I 13.40	I 13.60	I 13.51	I
SUBFRAME	I 3	I 11	I 19	I 27	I 35	I 43	I
COMPRESSION	I 5.14	I 4.74	I 5.22	I 4.74	I 4.23	I 5.50	I
MSE	I .575	I .625	I .401	I .185	I .303	I .222	I
OPS/PIXEL	I 13.39	I 13.35	I 13.43	I 13.40	I 13.36	I 13.46	I
SUBFRAME	I 2	I 10	I 18	I 26	I 34	I 42	I
COMPRESSION	I 3.22	I 3.36	I 6.48	I 4.15	I 4.52	I 6.89	I
MSE	I .825	I .196	I .751	I .272	I .310	I .304	I
OPS/PIXEL	I 13.16	I 13.23	I 13.47	I 13.35	I 13.38	I 13.53	I
SUBFRAME	I 1	I 9	I 17	I 25	I 33	I 41	I
COMPRESSION	I 2.78	I 4.31	I 5.88	I 5.33	I 6.18	I 5.89	I
MSE	I .665	I .168	I .798	I .439	I .196	I .600	I
OPS/PIXEL	I 13.10	I 13.33	I 13.46	I 13.43	I 13.45	I 13.47	I

Figure 4-3. Extreme Contrast Control Mode

(a) Compression/MSE/Operations Map




Figure 4-3(b). 48 Subframe Composite

STEP CONTRAST - LMAX=3

C= 4.97, MSE= .358, OPS=17.9, CONTROL 12 4 5 5
 8 3 4 4
 4 2 3 3
 0 0 0 0
 0 0 0 0
 0 0 0 0

SUBFRAME	I 8	I 16	I 24	I 32	I 40	I 48	I
COMPRESSION	I 2.84	I 10.33	I 5.27	I 8.55	I 4.65	I 12.88	I
MSE	I .313	I 1.407	I .213	I 1.324	I .817	I .244	I
OPS/PIXEL	I 17.13	I 18.41	I 17.93	I 18.31	I 17.82	I 18.47	I
SUBFRAME	I 7	I 15	I 23	I 31	I 39	I 47	I
COMPRESSION	I 3.88	I 8.59	I 5.65	I 7.14	I 8.92	I 4.02	I
MSE	I .891	I .891	I .131	I .761	I .555	I .336	I
OPS/PIXEL	I 17.60	I 18.32	I 17.95	I 18.21	I 18.31	I 17.51	I
SUBFRAME	I 6	I 14	I 22	I 30	I 38	I 46	I
COMPRESSION	I 6.04	I 4.47	I 7.99	I 6.16	I 13.18	I 3.55	I
MSE	I .779	I .730	I .148	I .478	I .220	I .401	I
OPS/PIXEL	I 18.09	I 17.77	I 18.19	I 18.05	I 18.46	I 17.32	I
SUBFRAME	I 5	I 13	I 21	I 29	I 37	I 45	I
COMPRESSION	I 6.10	I 3.93	I 7.87	I 4.54	I 11.80	I 1.31	I
MSE	I 1.007	I 1.066	I .134	I .197	I .154	I .207	I
OPS/PIXEL	I 18.08	I 17.64	I 18.19	I 17.74	I 18.40	I 15.19	I
SUBFRAME	I 4	I 12	I 20	I 28	I 36	I 44	I
COMPRESSION	I 5.82	I 3.44	I 5.55	I 5.15	I 23.74	I 6.57	I
MSE	I .939	I .680	I .154	I .426	I .111	I .253	I
OPS/PIXEL	I 18.04	I 17.44	I 17.98	I 17.92	I 18.59	I 18.09	I
SUBFRAME	I 3	I 11	I 19	I 27	I 35	I 43	I
COMPRESSION	I 5.14	I 4.86	I 5.63	I 5.01	I 3.87	I 5.04	I
MSE	I .620	I .698	I .468	I .215	I .294	I .216	I
OPS/PIXEL	I 17.94	I 17.86	I 18.03	I 17.88	I 17.61	I 17.87	I
SUBFRAME	I 2	I 10	I 18	I 26	I 34	I 42	I
COMPRESSION	I 3.25	I 3.46	I 7.07	I 4.31	I 4.33	I 6.38	I
MSE	I .912	I .227	I .864	I .312	I .322	I .306	I
OPS/PIXEL	I 17.37	I 17.40	I 18.20	I 17.74	I 17.73	I 18.09	I
SUBFRAME	I 1	I 9	I 17	I 25	I 33	I 41	I
COMPRESSION	I 2.83	I 4.46	I 6.29	I 5.14	I 6.23	I 5.33	I
MSE	I .748	I .191	I .907	I .453	I .214	I .589	I
OPS/PIXEL	I 17.18	I 17.71	I 18.13	I 17.92	I 18.03	I 17.94	I

Figure 4-4. Step Contrast Control Mode

(a) Compression/MSE/Operations Map




Figure 4-4(b). 48 Subframe Composite

needs two or three bits and the intensity code requires six bits. If all information had to be sent at the original resolution, the overhead of the resolution codes would result in a data explosion of one-third to one-half.

The block exhaustion constraint implies that for the special case of level zero (the original or highest resolution), elements always come in sets of four. Thus, the option exists to code each quadruplet of level zero pixels with a single two or three-bit resolution code (00 or 000) followed always by four intensity values. This limits the data explosion potential to between one-twelfth and one-eighth. Further, such 'quadruplet' coding (as opposed to 'singlet' coding where the resolution code is included with each level zero pixel) will raise the compression somewhat whenever level zero MAPS elements are present.

Quadruplet coding, however, destroys the uniformity of the MAPS data format since elements at level zero are coded as composites containing 26 or 27 bits while elements at level one or higher are coded as entities of 8 or 9 bits. A non-uniform data format in turn is much more vulnerable to degradation from transmission line errors. The latter show up as errors in value, not errors in occurrence. The question is not whether something was transmitted but whether it was a 'zero' or a 'one'. With a uniform data format, one knows where to look in the data stream for the next resolution code even if the current code is in error. This predictive capability disappears with a level-dependent format and a much more elaborate resynchronization test procedure is then required.

Two further observations assist in the resolution of this trade-off. First, target area compressions must be reasonably high if the overall compression goals are to be achieved. Thus, only a small portion of the target subframes can be retained at the level zero or original resolution. Second, since significantly reduced fidelity is acceptable in the background regions, the resolution therein will almost always be forced to at least level one or coarser values. Hence, the opportunities for compression enhancement using quadruplet level zero coding are severely limited in this high compression

application of MAPS. The selection of singlet level zero coding with its consequent data uniformity is a clear-cut choice given the expected error environment.

4.4 SELECTION OF MAPS MAXIMUM BLOCK SIZE

The maximum allowed MAPS block size defines the local subframe (smaller than and commensurate with the REARC subframe) which must be available to the process outlined previously in Figure 3-5. The 'level' of this largest block is denoted either as 'M' or 'LMAX'.

The maximum block size selection is complicated by three competing factors- the desire to minimize the size of the input buffer, the requirement to reach a high level of compression in the background regions (typically 40:1-60:1), and the discontinuous increase in overhead going from LMAX = 3 (8 x 8 blocks) to LMAX = 4 (16 x 16 blocks). This latter discontinuity arises because the element size codes for LMAX ≤ 3 can all be characterized with two bits per element but an increase to the range $4 \leq \text{LMAX} \leq 7$ forces use of three bits per element. Note, then, that several larger blocks ($L \leq 4$) must be formed just to overcome the additional overhead incurred by adding one more bit to every element. In general, each larger block is formed at the expense of some, albeit small, intensity differences separating the four component subblocks. Thus, the mean square error (MSE) for the image with larger blocks will be greater than that for the image with small blocks. In most instances, however, the larger blocks also increase the compression so something is gained. In the transition from LMAX = 3 to LMAX = 4, on the other hand, the resultant image may be worse in both MSE (larger) and compression (smaller). This conclusion applies over compression ranges where few of the largest blocks are employed; that is, where the potential of the maximum hasn't been or can't be exploited. This effect is illustrated in Figure 4-5 for several compressions of the forty-eight subframe composite image. The control matrix for $L \leq 3$ is the same for LMAX = 3 and LMAX = 4 for each point (x) plotted. This means that $\text{MSE (LMAX = 3)} \leq \text{MSE (LMAX = 4)}$ for every plotted point. The dashed line denotes equality of compression so it is seen that the compression reduction




Figure 4-5. Comparative Compressions for LMAX = 3 and LMAX = 4

due to the extra overhead bit is not overcome until compression ratios in the range of 15:1 to 20:1 are reached. Below this level, $L_{MAX} = 3$, is a clear-cut choice since both greater compression and smaller error result.

In this application, however, overall compressions in excess of 30:1 are required. With $L_{MAX} = 3$, a maximum possible compression of 48:1 results if all elements are forced to $L = L_{MAX}$ (i.e. 8×8 blocks). This figure arises from the ratio,

$$C_{MAX}(L_{MAX} = 3) = \frac{(8 \times 8 \text{ pixels}) (6 \text{ bits/pixel})}{(6 \text{ bits intensity} + 2 \text{ bits size}) \text{ per MAPS element}} = 48:1. (4.2)$$

This maximum value (which is shown in Figure 4-5) is too small to reach the required overall compression since toward the high end (over about 25:1) there is not enough room to 'maneuver' and too many blocks must be forced to $L = 3$. In effect, the contrast thresholds must be set too high and the natural spatial variation of the image cannot be exploited.

At $L_{MAX} = 4$, the maximum compression ratio is:

$$C_{MAX}(L_{MAX} = 4) = \frac{(16 \times 16 \text{ pixels}) (6 \text{ bits/pixel})}{(6 \text{ bits intensity} + \underline{3} \text{ bits size}) \text{ per MAPS element}} = 170.7:1. (4.3)$$

Similar calculations yield $C_{MAX}(L_{MAX} = 5) = 682.7:1$ and $C_{MAX}(L_{MAX} = 6) = 2730.7:1$. In practice, it is found that few images possess the spatial dynamic range to make effective use of the (32×32) and (64×64) blocks in more than token numbers. Thus, these levels add input buffer memory (doubling with each step in L_{MAX}) while providing only slight gains in compression; the latter less than 10% at the 60:1 level. Thus, the selection $L_{MAX} = 4$ appears to provide:

- adequate dynamic range to reach 30:1 compression,
- minimum input buffer memory consistent with this compression,
- consistency with the useful spatial dynamic range of typical imagery.

On this basis, $L_{MAX} = 4$ is selected as the MAPS subframe size.

Note that a further alternative would involve using $L_{MAX} = 4$ in the background and $L_{MAX} = 3$ in the target regions. This could boost compression (at a given error) but two significant penalties are incurred. First, the data loses its 'transparency' to target/background specifications and some additional overhead results to characterize the transitions back and forth. The overhead is not severe with an 8×8 subframe specification such as REARC but it would significantly complicate the more general interactive macro-fidelity variation capability toward which MAPS is evolving. The second penalty is the same as that which decided the $L = 0$ singlet vs quadruplet coding trade -- viz., data element uniformity. Using $L_{MAX} = 4$ throughout the image means that each MAPS element consists of a constant length bit string (9 bits) so the location of the next size code in the data stream can be predicted unambiguously regardless of possible bit errors. This vastly simplifies error handling which is crucial in this application.

4.5 SUMMARY OF INTRA-MAPS TRADE-OFF SELECTIONS

The three intra-MAPS trade-offs and their general impacts on the key performance variables are summarized in Figure 4-6. The final trade-offs are resolved as shown at the bottom of this chart. These selections are used in all of the full-frame MAPS compression examples which are presented in the next section.

PERFORMANCE	TRADE-OFF		
	CONTRAST CONTROL MODE	HIGHEST RESOLUTION CODE	MAXIMUM BLOCK SIZE
COMPRESSION		FAVORS QUADRUPLER	FAVORS: 8x8 IN TARGET 64x64 IN BACKGROUND
FIDELITY	FAVORS STEP		
ERROR RECOVERY		FAVORS SINGLET	FAVORS CONSTANT
OPERATIONS COUNT	FAVORS EXTREME		
MEMORY			FAVORS SMALLER
<u>SELECTION</u>	<u>STEP</u>	<u>SINGLET</u>	<u>16 x 16</u>

Figure 4-6. Intra-MAPS Trade-off Selections

SECTION FIVE

TEST IMAGERY RESULTS

Ten diverse full frame (2048 x 2048) images constitute the test set for the current application and eight of these contain target/background partitions. The target regions vary from four to eight REARC subframes (256 x 256) per image.

This section presents the performance of MAPS on the test set. In addition, a model for transmission line errors is described and the effects of these errors are illustrated.

5.1 TABULATION OF MAPS PERFORMANCE

Micro-Adaptive Picture Sequencing (MAPS) controls image fidelity (in the form of local contrast) directly. Hence, the resultant compression is a derived quantity which depends jointly on the MAPS contrast threshold matrices and the contents of the specific image; the compression ratio is not known a priori.

Since compression is not directly selectable, several MAPS runs were made on each image and an example for which the overall compression just exceeds 30:1 was retained for each scene. This yields a common comparison set for the purposes of the development tests. In practice, the desired fidelity would be selected and the compression would be allowed to vary from image to image according to their respective contents. In this way, the variable 'compressibility' exhibited by typical imagery can be exploited without forcing excessive degradation on some scenes and not taking advantage of the compression potential on others.

In addition to the set of ten '30:1' compressions, five of the scenes were subjected to two further compressions each to illustrate the trade-off between compression ratio and fidelity. In each of these cases, the 30:1

result was used as a guide. Where significant degradation was seen at 30:1, the two alternate runs were made at lower compression. Otherwise, the two alternatives were chosen to bracket the 30:1 case.

MAPS performance for these twenty examples is summarized in Tables 5-1, 5-2, and 5-3. Each of these tabulations is keyed by the scene title and the overall compression ratio. This provides the cross reference among the tables and among the hardcopy imagery of the next subsection.

Table 5-1 lists the numbers of REARC target subframes and the compression ratios for the target and background regions separately.

Table 5-2 presents the fidelity performance for target and background in terms of the mean square error (see Equation 4.1). All MSE values are in percent. Note in some instances that the MSE in the target area is actually worse than that in the background in spite of the fact that the target compressions are much lower. In part this is a reflection of the 'busyness' of the target regions. It is also attributable to a small denominator term in Equation 4.1 for the target region of the Aerial Photo scene; this results in an artificially high MSE value.

Finally, Table 5-3 characterizes the overall operations counts for the full scenes. Note that this count varies by only about 2% while the compression ratio varies by a factor of 3.85.

5.2 MAPS IMAGERY

Ten original images and their corresponding twenty MAPS decompressions are presented as Figures 5-1(a) through 5-10(b). The originals show the REARC target regions outlined with a border consisting of one white and one black pixel.

All MAPS examples were decompressed using the very fast block decompression mode. Examples of the current version of the adaptive decompression mode can be seen in Figures 3-7, 3-8, and 3-9. This latter technique is still evolving as noted in Section Seven.

TABLE 5-1. FULL FRAME COMPRESSIONS - COMPRESSION RATIO

FULL FRAME COMPRESSIONS				
IMAGE	TARGET BLOCKS	COMPRESSION RATIO		
		TARGET	BACKGROUND	OVERALL
AERIAL PHOTO	6	2.86	24.89	14.46
		4.38	37.15	21.84
		7.38	44.62	30.29
RIVERFRONT	4	9.09	35.46	30.02
RADAR SCENE	4	5.85	24.68	20.55
		8.58	36.25	30.17
		11.82	45.98	38.95
HARBOR 8 x 8	4	6.22	41.77	30.78
HARBOR 4 x 4	8	9.98	27.61	22.62
		11.31	39.34	30.03
		12.87	53.32	38.28
HARBOR 2 x 2	4	5.40	43.13	30.03
IEEE CHART	64	-	-	10.12
		-	-	16.26
		-	-	30.76
SAM SITE	6	9.02	39.70	30.10
AIR FIELD	8	8.31	28.43	21.83
		9.52	44.13	30.34
		9.52	57.70	35.34
IEEE GIRL	64	-	-	30.54

TABLE 5-2. FULL FRAME COMPRESSIONS - MEAN SQUARE ERROR

<u>IMAGE</u>	<u>OVERALL COMPRESSION</u>	<u>MSE (MEAN SQUARE ERROR)</u>	
		<u>TARGET</u>	<u>BACKGROUND</u>
AERIAL PHOTO	14.46	0.620	0.718
	21.84	0.921	0.861
	30.29	1.363	0.965
RIVERFRONT	30.02	1.035	3.340
RADAR SCENE	20.55	1.144	2.193
	30.17	1.438	2.534
	38.95	1.736	2.792
HARBOR 8 x 8	30.78	0.495	0.554
HARBOR 4 x 4	22.62	0.360	0.518
	30.03	0.403	0.663
	38.28	0.445	0.774
HARBOR 2 x 2	30.03	0.316	0.366
IEEE CHART	10.12	0.822	-
	16.26	1.341	-
	30.76	2.646	-
SAM SITE	30.10	0.655	0.548
AIRFIELD	21.83	0.407	0.488
	30.34	0.478	0.684
	35.34	0.478	0.819
IEEE GIRL	30.54	0.175	-

TABLE 5-3. FULL FRAME COMPRESSIONS - OPERATIONS COUNT

FULL FRAME COMPRESSIONS		
<u>IMAGE</u>	<u>OVERALL COMPRESSION</u>	<u>OPERATIONS PER PIXEL</u>
AERIAL PHOTO	14.46	18.84
	21.84	18.94
	30.29	19.00
RIVERFRONT	30.02	18.95
RADAR SCENE	20.55	18.93
	30.17	18.99
	38.95	19.03
HARBOR 8 x 8	30.78	18.93
HARBOR 4 x 4	22.62	18.89
	30.03	18.93
	38.28	18.96
HARBOR 2 x 2	30.03	18.90
IEEE CHART	10.12	18.63
	16.26	18.80
	30.76	18.93
SAM SITE	30.10	18.97
AIRFIELD	21.83	18.89
	30.34	18.93
	35.34	18.95
IEEE GIRL	30.54	18.88




Figure 5-1(a). Aerial Photo - Targeted Original




Figure 5-1(b). Aerial Photo - 14.46:1




Figure 5-1(c). Aerial Photo - 21.84:1




Figure 5-1(d). Aerial Photo - 30.29:1



Figure 5-2(a). Riverfront - Targeted Original



Figure 5-2(b). Riverfront - 30.02:1



Figure 5-3(a). Radar Scene - Targeted Original




Figure 5-3(b). Radar Scene - 20.55:1



Figure 5-3(c). Radar Scene - 30.17:1




Figure 5-3(d). Radar Scene - 38.95:1




Figure 5-4(a). Harbor 8 x 8 - Targeted Original



Figure 5-4(b). Harbor 8 x 8 - 30.78:1



Figure 5-5(a). Harbor 4 x 4 - Targeted Original



Figure 5-5(b). Harbor 4 x 4 - 22.62:1



Figure 5-5(c). Harbor 4 x 4 - 30.03:1



Figure 5-5(d). Harbor 4 x 4 - 38.28:1




Figure 5-6(a). Harbor 2 x 2 - Targeted Original




Figure 5-6(b). Harbor 2 x 2 - 30/03:1




Figure 5-7(b). IEEE Chart - 10.12:1



Figure 5-7(c). IEEE Chart - 16.26:1




Figure 5-7(d). IEEE Chart - 30.76:1




Figure 5-8(a). SAM Site - Targeted Original




Figure 5-8(b). SAM Site - 30.10:1




Figure 5-9(a). Airfield - Targeted Original




Figure 5-9(b). Airfield - 21.83:1




Figure 5-9(c). Airfield - 30.34:1




Figure 5-9(d). Airfield - 35.34:1



Figure 5-10(a). IEEE Girl - Original



Figure 5-10(b). IEEE Girl - 30.54:1

5.3 LINE ERROR EFFECTS

Transmission over land lines is subject to significant bit error rates with values as high as 10^{-3} (9600 baud line). Such errors affect the transmission of MAPS-encoded imagery in two distinct ways. If an error occurs in the intensity portion of a MAPS element, the effect is simply to 'paint' the corresponding block with the wrong gray scale. This will be noticeable only if (1) the block is fairly large and (2) the error occurs in a high order intensity bit. Even under these conditions, the effects of the error remain localized and affect no more than a 16×16 block (since $IMAX = 4$). Note that such intensity errors are not directly detectable in the data stream without the addition of explicit error checking bits; they show up only through their effect on the visual image. Fortunately, the source of any noticeable intensity errors is immediately obvious and they would rarely affect the information content of the image. Thus, it is not recommended that the compression be reduced by providing the bits necessary for intensity error checking.

The potential effects of the second type of error - changed bits in the resolution portion of a MAPS element - are much more severe. Such errors will quickly cause the implicit position sequence to be destroyed. It is generally possible to recover from errors in the resolution codes, however, because they retain a very useful amount of redundancy. Indeed, this recovery potential is inherent in the raw MAPS bit stream and is available without additional code bits; the compression is not affected.

In essence, the MAPS sequence convention and the block exhaustion constraint (Section 3.1.1.1) imply a set of consistency relations which the pattern of resolution codes must meet. The MAPS elements must fit together like a jigsaw puzzle!

Compound errors may occasionally leave ambiguities in the resolution of a detected error pattern. However, recovery in these pathological cases can be effected by a simple coding artifice which again is available with no added bits. The resolution portion of the MAPS element consists of three bits but

with the selected maximum level, $L_{MAX} = 4$, only five of the possible eight codes are used. Thus, the last resolution code in every 16×16 MAPS subframe partition is replaced by one of these unused codes. The true value of the code for the final element is implicit (assuming no errors) and the special terminator provides both a higher level check point and a restart location on a MAPS-subframe-by-subframe basis.

Optimal implementation of this MAPS error recovery potential is another area of evolution for the technique. The nature of the errors anticipated is presented in the next subsection.

5.3.1 Line Error Model

The distribution of bit errors on a transmission line does not follow any of the more familiar random patterns since such errors tend to occur in "bursts" separated by varying-length periods of quiescence. The line error simulation chosen here is based on the truncated Pareto distribution. This approach gives rise to a stream of inter-error intervals which shows the same 'burst' statistics as are observed on typical land lines and the entire distribution has been seen to apply to phone lines over sequences at least out to 10^6 bits⁶.

Since a typical 2048×2048 image with a 30:1 overall compression ratio contains about 0.84×10^6 bits, this model appears to be fully appropriate.

The relevant relations for the model are summarized in Figure 5-11. Two parameters, T and α completely characterize the distribution. Here T is the truncation time which corresponds to the longest allowed inter-error spacing and α is adjusted (given T) to obtain the desired average error rate $\langle t \rangle^{-1}$. Curves of the mean error rate versus α with T parametric are presented in Figure 5-12. For a truncation point corresponding to 10^6 'bit times', it is seen that a value $\alpha = 0.5$ is required to yield a 10^{-3} mean error rate. This was chosen as the operating state for the current simulation.

LINE ERROR SIMULATION

TRUNCATED PARETO DISTRIBUTION

CUMULATIVE DISTRIBUTION:

$$f(t) = \frac{1 - t^{-\alpha}}{1 - T^{-\alpha}}, \quad 1 \leq t \leq T$$

DENSITY:

$$\frac{df(t)}{dt} = \frac{\alpha t^{-(1+\alpha)}}{1 - T^{-\alpha}}$$

MEAN INTERERROR INTERVAL:

$$\langle t \rangle = \frac{\alpha}{1-\alpha} \frac{T^{1-\alpha} - 1}{1 - T^{-\alpha}}$$

NEXT INTERERROR INTERVAL SELECTION:

$$t = [1 - R(1 - T^{-\alpha})]^{-1/\alpha}$$

$$\left. \begin{array}{l} T = 10^{+6} \\ \alpha = 0.5 \end{array} \right\} \rightarrow \langle t \rangle = 10^3$$

Figure 5-11. Line Error Model Relations

AD-A050 679

CONTROL DATA CORP MINNEAPOLIS MINN DIGITAL IMAGE SYS--ETC F/G 14/5
IMAGE COMPRESSION TECHNIQUES.(U)

DEC 77 A E LABONTE, C J MCCALLUM

F30602-76-C-0350

UNCLASSIFIED

RADC-TR-77-405

NL

2 OF 2

AD
A050679



END

DATE

FILMED

4 -78

DDC

Mean Error Rate

TRUNCATED PARETO DISTRIBUTION




Figure 5-12. Mean Error Rate vs α for the Truncated Pareto Distribution




Figure 5-13. Line Error Intensity Effects: Airfield - 30.34:1

SECTION SIX

IMPLEMENTATION RECOMMENDATIONS

Although several areas for further evolution of MAPS are described in Section Seven, a complete basic version of the technique could be implemented immediately based on the concept already described. It is recommended that this initial implementation consist of microcode for an existing machine such as the Control Data Corporation[®] Flexible Processor which is specifically tailored for image processing. This approach preserves the capability to incorporate evolving MAPS enhancements through software modification.

The CDC[®] Flexible Processor is micro-programmable with up to 4K of 48-bit micromemory and up to 2K of 32-bit register file storage. The companion image memory can be configured up to two megabytes. Execution speed is 8×10^6 instructions per second or about 30×10^6 elementary operations per second. These operating characteristics insure sufficient capacity to embody the expected growth of MAPS.

The basic physical characteristics of the flexible processor are 500 watts of power consumption, approximately one cubic foot of volume in a 19" rack mount configuration, and weight in the 30 to 40 lb. range. Cost of the basic flexible processor falls in the \$30K to \$40K bracket with a typical image memory (131 kilobytes) at about half of this. The exact specifications vary slightly depending on the particular configuration selected.

Software implementation of MAPS is straightforward based on the procedures documented in Sections 3.1.1 and 3.1.2.1 with the basic flow charted in Figure 3-5. Two observations beyond those of Section Three are helpful in obtaining an efficient implementation. First, the determination of the contrasts is essentially a four-element sort problem. The number of elements is sufficiently small so that a full tree sort is both practical and optimal from a speed standpoint. Second, since a maximum block size of 16×16 is selected, each MAPS subframe contains only 256 original pixels. Thus, a table lookup to implement the addressing sequence in Figure 3-2c is recommended.

SECTION SEVEN

MAPS EVOLUTION

Immediate MAPS performance enhancement would accrue from developments in five task areas - adaptive decompression, control matrix selection, error recovery optimization, interactive macro-fidelity control, and on-line MAPS demonstration. Brief outlines of each area are given in the following subsections.

7.1 ADAPTIVE DECOMPRESSION

The potential for subjective image improvement through adaptive decompression goes well beyond the preliminary results given in Section 3.1.2.2. MAPS exhibits a natural compatibility with adaptive decompression since the 'activity' level in an area of the scene is keyed directly by the MAPS resolution code. Thus, a prior measure of the local correlation length is immediately available.

This a priori knowledge may then be used to select both the size of the convolution window and the elements within that window which should be used. This latter selection is based on both element size and intensity relative to the size and intensity of the central element. In addition, several options for the convolution weighting function - e.g. uniform, pyramid, or $\sin x/x$ - should be explored. Finally, the effects of 'contouring' between large uniform adjacent regions of slightly different gray scale can be suppressed by adding a 'dither' to the decompression operation. Again, the amplitude of this dither can be made adaptive to the prior knowledge contained in the MAPS resolution codes.

Detailed trade-off studies along these several dimensions of adaptive decompression are recommended as the first task for further MAPS development. Throughout these studies, the additional dimension of 'algorithm optimization' for cost and speed should also be addressed.

7.2 CONTROL MATRIX SELECTION

As the second task, guidelines for selection of the MAPS contrast control matrices should be extended to give a rapid classification scheme for production applications of the technique. The classes should include:

- Type of imagery (photo, radar, line drawing, text, etc.),
- Characteristics of imagery (contrast, scale, etc.), and
- Content of imagery (cultural features, natural features).

A simple check list or decision tree should be constructed to aid operational selection. Override capability to 'fine-tune' the matrix with operator-selected values should also be retained.

For the longer term, automatic control selection based on global sampling of the image statistics should be investigated. Here, a suitable metric relating these statistics and the effects of the control matrix must be developed.

7.3 ERROR RECOVERY OPTIMIZATION

At the selected MAPS subframe size of 16×16 ($L_{MAX} = 4$), there are just under 5×10^{19} distinct legitimate geometric configurations of MAPS elements! Thus, an exhaustive examination of the possible pathological error effects in the MAPS resolution codes is out of the question.

However, the three unused resolution codes (out of eight possible) can be utilized in a more optimal fashion than the assignment of a single sub-frame termination code as outlined in Section 5.3. Exploration of these options and development of the corresponding optimized error detection/recovery algorithm is recommended as the third near-term task in the evolution of MAPS.

7.4 INTERACTIVE MACRO-FIDELITY CONTROL

MAPS presents the opportunity to replace the variable fidelity

selection embodied in the REARC technique with an even more powerful macro-fidelity control capability. In theory, each MAPS 16 x 16 subframe could be given its own contrast control matrix with fidelity ranging from perfect preservation to complete obliteration of all intra-subframe detail. Thus, the output image could be tailored to retain just those features which are relevant and at just the minimum local resolution required.

The immense power of MAPS in this case resides in the fact that all 'scaffolding' used to characterize the variable fidelity during the coding step is discarded once the MAPS encoding is complete. Thus, transmission, storage, and compression are 'transparent' to these controls; all relevant MAPS information is carried locally in the individual resolution/intensity elements.

Since the MAPS subframe is small, regions of almost arbitrary boundary can be circumscribed on an interactive display and then MAPS-compressed to the minimum acceptable fidelity (maximum compression). If necessary, subregions within the originally defined region can then be designated and individual features restored or deleted as desired.

Development and demonstration of this macro-fidelity control capability is the fourth recommended task.

7.5 ON-LINE MAPS DEMONSTRATION

As the final near-term evolution step, implementation of MAPS in an on-line demonstration mode is recommended. Here, the system might be incorporated in the Experimental Image Compression Subsystem (EICS) at RADC with the possible requirement for addition of a CDC[®] Flexible Processor and interface dependent on a detailed evaluation of available computation and memory capacity. Alternatively or perhaps coincidentally, the system might be incorporated in microcode in the existing interactive CDC[®] Cyber-Ikon image processing facility at Control Data Corporation.

Such implementation will provide an on-going test-bed for empirical exploration and evolution of the robust potential contained in the basic concepts of Micro-Adaptive Picture Sequencing.

References

1. Sillers, W. M., "Experimental Image Compression Subsystem (EICS)", Proceedings of the Society of Photo-Optical Instrumentation Engineers, Volume 87, Advances in Image Transmission Techniques, pp 47-54, 1976.

The following three documents give extensive collections of articles on both classical and current developments in image compression.

2. Jayant, N. S. (editor), WAVEFORM QUANTIZATION AND CODING, IEEE Press Selected Reprint Series, New York, 1976.
3. Proceedings of the Society of Photo-Optical Instrumentation Engineers, Volume 66, Efficient Transmission of Pictorial Information, 1975.
4. Proceedings of the Society of Photo-Optical Instrumentation Engineers, Volume 87, Advances in Image Transmission Techniques, 1976.
5. Ahmed, N. and K. Rao, ORTHOGONAL TRANSFORMS FOR DIGITAL SIGNAL PROCESSING, Sections 9.4 and 9.5 Springer-Verlog, 1975.
6. Sussman, S. M., "Analysis of the Pareto Model for Error Statistics on Telephone Circuits," IEEE Transactions on Communication Systems, pp 213-221, June, 1963.